

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
1 December 2005 (01.12.2005)

PCT

(10) International Publication Number
WO 2005/115026 A2

(51) International Patent Classification⁷: **H04Q 7/20**

(21) International Application Number:
PCT/US2005/017385

(22) International Filing Date: 17 May 2005 (17.05.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/848,597 18 May 2004 (18.05.2004) US
11/037,896 18 January 2005 (18.01.2005) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:

US 11/037,896 (CON)
Filed on 18 January 2005 (18.01.2005)
US 10/848,597 (CIP)
Filed on 18 May 2004 (18.05.2004)

(71) Applicant (for all designated States except US): **AIR-VANA, INC.** [US/US]; 25 Industrial Avenue, Chelmsford, MA 01720 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **EYUBOGLU,**

Vedat [US/US]; 150 Jennie Dugan Road, Concord, MA 01742 (US). **KIM, Woojune** [KR/US]; 53 Lake St #2, Arlington, MA 02474 (US). **BARABELL, Arthur, J.** [US/US]; 11 Hayden Circle, Sudbury, MA 01776-2098 (US). **CHERIAN, Sanjay** [US/US]; 6 Maxwell Drive, Brookline, MA 03033 (US).


(74) Agent: **FEIGENBAUM, David, L.**; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US (patent), UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

[Continued on next page]

(54) Title: RADIO NETWORK CONTROL



(57) Abstract: A first session is established for a first mobile access terminal on a first radio network controller via a first radio node. A first traffic channel is established between the first mobile access terminal and the first radio network controller. A first plurality of packets are sent and received over the first traffic channel. The first plurality of packets travel between a first radio node and the first radio network controller without passing through a second radio network controller. The first traffic channel is maintained as the first access terminal moves from a coverage area of the first radio node to a coverage area of a second radio node. A second plurality of packets travel between the second radio node and the first radio network controller without passing through another radio network controller. In a radio access network, multiple radio network controllers are connected to several radio nodes using a network. The interconnected radio network controllers and radio nodes are addressable, and, therefore, each radio network controller can communicate directly with each radio node and visa versa. The radio access network can be configured to avoid active handoffs between radio network controllers by maintaining a traffic channel set up between an access terminal and a radio network controller even as the flow.



WO 2005/115026 A2



FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— *without international search report and to be republished upon receipt of that report*

ABSTRACT

A first session is established for a first mobile access terminal on a first radio network controller via a first radio node. A first traffic channel is established between the first mobile access terminal and the first radio network controller. A first plurality
5 of packets are sent and received over the first traffic channel. The first plurality of packets travel between a first radio node and the first radio network controller without passing through a second radio network controller. The first traffic channel is maintained as the first access terminal moves from a coverage area of the first radio node to a coverage area of a second radio node. A second plurality of packets travel
10 between the second radio node and the first radio network controller without passing through another radio network controller. In a radio access network, multiple radio network controllers are connected to several radio nodes using a network. The interconnected radio network controllers and radio nodes are addressable, and, therefore, each radio network controller can communicate directly with each radio node
15 and visa versa. The radio access network can be configured to avoid active handoffs between radio network controllers by maintaining a traffic channel set up between an access terminal and a radio network controller even as the flow.

21089123.doc

RADIO NETWORK CONTROL

BACKGROUND

This description relates to radio network control.

High Data Rate (HDR) is an emerging mobile wireless access technology that
5 enables personal broadband Internet services to be accessed anywhere, anytime (see P.
Bender, et al., "CDMA/HDR: A Bandwidth-Efficient High-Speed Wireless Data
Service for Nomadic Users", IEEE Communications Magazine, July 2000, and 3GPP2,
"Draft Baseline Text for 1xEV-DO," August 21, 2000). Developed by Qualcomm,
HDR is an air interface optimized for IP packet data services that can deliver a shared
10 forward link transmission rate of up to 2.46 Mbit/s per sector using only (1X) 1.25
MHz of spectrum. Compatible with CDMA2000 radio access (TIA/EIA/IS-2001,
"Interoperability Specification (IOS) for CDMA2000 Network Access Interfaces," May
2000) and wireless IP network interfaces (TIA/EIA/TSB-115, "Wireless IP
Architecture Based on IETF Protocols," June 6, 2000, and TIA/EIA/IS-835, "Wireless
15 IP Network Standard," 3rd Generation Partnership Project 2 (3GPP2), Version 1.0,
July 14, 2000), HDR networks can be built entirely on IP technologies, all the way
from the mobile Access Terminal (AT) to the global Internet, thus taking full advantage
of the scalability, redundancy and low-cost of IP networks.

HDR has been adopted by TIA (Telecommunications Industry Association) as a
20 new standard in the CDMA2000 family, an EVolution of the current 1xRTT standard
for high-speed data-only (DO) services, formally referred to as HRPD (High Rate
Packet Data), also known as 1xEV-DO or IS-856.

IS-856 systems are typically implemented using the radio access network
architecture shown in Figure 1. Here the Access Terminal (AT) 10 may be a laptop
25 computer, a Personal Digital Assistant (PDA), a dual-mode voice/data handset, or
another device, with built-in IS-856 support.

The entire administrative service area of a wireless access provider may be
divided into one or more subnetworks (or subnets) 12, 14. Each subnet 12 includes a
set of Radio Nodes (RN's) 16, 18 and one or more Radio Network Controllers (RNC)
30 20, 22. The RN's are connected to RNC's over a backhaul network 24. In existing 2G
and 3G wireless networks, each RN is connected to only one RNC using dedicated

leased lines or ATM permanent virtual circuits (PVC's). Further, RNC's are connected to each other using dedicated leased lines or ATM PVC's. In a new generation of IP-based radio access networks, the backhaul can be implemented using a shared IP or metropolitan Ethernet network which supports many-to-many connectivity between
5 RN's and RNC's.

Each RNC typically controls 25-100 RN's. Each RN typically supports 1-4 carriers each of 1.25 MHz of bandwidth. Further, each cell area (not shown) is typically divided into multiple sectors (typically 3 or 6) and the RN has one radio transceiver 27 for each sector.

10 Each RNC is connected over an IP network 26 to one or more Packet Data Serving Node's (PDSN's) 28 (see TIA references cited earlier). The RNC communicates with the PDSN over a standard interface termed the R-P (Radio-Packet) interface 30. The R-P interface is further broken into two interfaces: the A10 interface used to carry data and the A11 interface used to carry signaling. A PDSN can be
15 viewed as an edge router that supports mobility; it maintains link layer connectivity to AT's through the Access Network. The PDSN also interfaces to AAA servers 32 for Authentication, Authorization, and Accounting (AAA).

Once an AT connects to the network, it establishes session with an RNC and receives a link layer address from the RNC. The session represents all the information
20 the RNC needs to serve the AT. In IS-856 radio access networks as currently defined by 3GPP2 in 1xEV-DO IOS Phase 1 (IS-878), each RN is uniquely associated with an RNC and each subnet contains only one RNC. As a result, when an AT moves from the coverage area of one RNC to the coverage area of another, the AT performs a handoff, which includes a session transfer.

25 Every time a dormant AT crosses a subnet boundary, the AT initiates a dormant handoff by sending a UATI_Request. The AT recognizes the need for a dormant handoff by monitoring the 128-bit SectorID being broadcast by the sectors. All sectors that belong to the same subnet have SectorID's that fall within a certain range. The 128-bit Universal Access Terminal Identifier (UATI) assigned to an AT in a given
30 subnet falls within the same range. When the AT moves into the coverage area of another subnet, the AT compares its UATI with the SectorID being broadcast by its serving sector. When these do not belong to the same range, the AT knows that it has

crossed a subnet boundary and initiates the dormant handoff by sending a UATI_Request.

A first purpose of a dormant handoff is to inform the PDSN to send packets arriving for that AT to the new serving RNC. Dormant handoffs involve a relocation of the R-P (A10) session from the old serving RNC to the new serving RNC. Without
5 such handoffs, the PDSN would send packets to an old serving RNC. Since the old serving RNC does not know the location of the AT outside its subnet, AT's packets may be lost.

A second purpose of a dormant handoff is to transfer session information
10 between RNC's. In IS-856, each RNC maintains certain session information about the AT. Such session information is needed for communication over the air interface. Session information includes the Universal Access Terminal Identifier (UATI), security keys for access channel authentication and encryption, and other protocol constants. Every time the AT crosses an RNC boundary (in this case a subnet), a new UATI needs
15 to be assigned to the AT and the remaining session information needs to be transferred from the old serving RNC to the new serving RNC. Such a transfer requires a network link between the RNC's. Without such session transfer, every handoff between RNC's would result in a new and lengthy session establishment, taking up precious air resources and causing delays. When the footprint of an RNC is small, dormant
20 handoffs occur frequently, resulting in excessive use of airlink resources (for the new UATI assignment), extra processing for the RNC's to implement the session transfer, and extra processing for the RNC and PDSN to relocate the A10 connection.

SUMMARY

In one aspect, there is a method. The method includes, in connection with a
25 mobile wireless network including a first and a second radio network controller and a first and a second radio node, establishing a first session for a first mobile access terminal on the first radio network controller via the first radio node. The method also includes establishing a second session for a second mobile access terminal on the second radio network controller via the second radio node and establishing a first traffic
30 channel between the first mobile access terminal and the first radio network controller. The method also includes sending and receiving a first plurality of packets over the first traffic channel, where the first plurality of packets traveling between a first radio node

and the first radio network controller without passing through the second radio network controller. The method also includes maintaining the first traffic channel as the first access terminal moves from a coverage area of the first radio node to a coverage area of the second radio node, where a second plurality of packets travel between the second
5 radio node and the first radio network controller without passing through another radio network controller. The method also includes establishing a new session for the first mobile access terminal on the second radio network controller or performing an inter-RNC dormant handoff from the first radio network controller to the second radio network controller when the first access terminal is in a dormant state after moving
10 from a coverage area of the first radio node to a coverage area of the second radio node.

In another aspect, there is a method. The method includes, in connection with a mobile wireless network including a first and a second radio network controller and a first and a second radio node, the radio network controllers each implemented on a chassis-based hardware platform with multiple server cards, establishing a first session
15 for a first mobile access terminal on the first radio network controller via the first radio node. The method also includes establishing a first traffic channel between a first mobile access terminal and the first radio network controller and sending and receiving a first plurality of packets over the first traffic channel, where the first plurality of packets travel between a first radio node and the first radio network controller without
20 passing through any other radio network controller. The method also includes establishing a second traffic channel between a second mobile access terminal and the second radio network controller and sending and receiving a second plurality of packets over the second traffic channel, where the second plurality of packets travel between the second radio node and the second radio network controller without passing through
25 any other radio network controller. The method also includes maintaining the first traffic channel as the first access terminal moves from a coverage area of the first radio node to a coverage area of the second radio node, a third plurality of packets traveling between the second radio node and the first radio network controller without passing through another radio network controller, and maintaining the first session when the
30 first mobile access terminal moves to the coverage area of the second radio node.

In another aspect, there is a computer program product, tangibly embodied in an information carrier, and adapted to work in a mobile wireless network including a first

and a second radio network controller and a first and a second radio node. The computer program product includes instructions operable to cause data processing apparatus to establish a first session for a first mobile access terminal on the first radio network controller via the first radio node and establish a second session for a second
5 mobile access terminal on the second radio network controller via the second radio node. The computer program product also includes instructions that are further operable to cause the data processing apparatus to establish a first traffic channel between the first mobile access terminal and the first radio network controller and send and receive a first plurality of packets over the first traffic channel, where the first
10 plurality of packets travel between a first radio node and the first radio network controller without passing through the second radio network controller. The computer program product also includes instructions that are further operable to cause the data processing apparatus to maintain the first traffic channel as the first access terminal moves from a coverage area of the first radio node to a coverage area of the second
15 radio node, where a second plurality of packets travel between the second radio node and the first radio network controller without passing through another radio network controller.

In another aspect, there is a mobile wireless network. The mobile wireless network includes a first radio network controller, a second radio network controller, a
20 first radio node, a second radio node, a first mobile access terminal, and a second mobile access terminal. The first mobile access terminal is associated with a first session established on the first radio network controller via the first radio node and a first traffic channel established with the first radio network controller. The first mobile access terminal sends and receives a first plurality of packets over the first traffic
25 channel, wherein the first plurality of packets travel between a first radio node and the first radio network controller without passing through the second radio network controller. The second mobile access terminal is associated with a second session established on the second radio network controller via the second radio node. The first traffic channel is maintained as the first access terminal moves from a coverage area of
30 the first radio node to a coverage area of the second radio node. A second plurality of packets travel between the second radio node and the first radio network controller without passing through another radio network controller.

Other examples of any of the aspects above can include one or more of the following features. A primary association can be established between the second radio node and the second radio network controller. A secondary association can be established between the second radio node and the first radio network controller.

- 5 Establishing the primary association can include passing PN offset and IP address information from the second radio node to the second radio network controller. Establishing the secondary association can include passing PN offset and IP address information from the second radio node to the first radio network controller. The first radio node can broadcast a first subnet identifier. The second radio node can broadcast
10 a second subnet identifier different from the first subnet identifier. Each radio node can be individually configured with its respective subnet identifier. The radio nodes can each obtain its subnet identifier from its respective radio network controller with whom it has established a primary association.

The first access terminal can monitor the subnet identifier while in the dormant
15 state and trigger the new session establishment or the dormant inter-RNC handoff from the first radio network controller to the second radio network controller upon detecting a change in the subnet identifier. The first access terminal can trigger the new session establishment or the dormant inter-RNC handoff from the first radio network controller to the second radio network controller by sending a UATI Request message.

20 A primary association can be established between a third radio node and the first radio network controller. A secondary association can be established between the third radio node and the second radio network controller. The third radio node can broadcast a third subnet identifier different from first and second subnet identifiers associated with the first and second radio nodes. The first radio network controller can
25 assign a new UATI to the first access terminal when it is in the dormant state after moving from a coverage area of the first radio node to a coverage area of the third radio node. An RNC resource control agent can be employed for storing the session information for sessions being served by one or more radio network controllers. The RNC resource control agent can detect a failure of a radio network controller and, upon
30 detection of the failure of a radio network controller, reassign user sessions to remaining radio network controllers and pass session information to these remaining radio network controllers. Partial session information can be passed that provides a

new radio network controller enough information about the first access terminal to send a message to close the session.

A chassis-based hardware platform with multiple server cards can be employed to implement each of the radio network controllers. An association can be established
5 between the second radio node and the first radio network controller by homing on one of the server cards on the first radio network controller. An association can be established between the second radio node and the second radio network controller by homing on one of the server cards on the second radio network controller. A transport-layer connection, such as TCP or SCTP, can be established for signaling between each
10 radio node and the server cards to which it is homed in the radio network controllers. Establishing the association can include passing PN offset information from the second radio node to the server card to which it is homed. The PN offset information can be distributed from the server cards receiving the PN offset information to other server cards in the chassis-based hardware platform. The inter-RNC handoff procedure can
15 comply with an A13 interface of a 1xEV-DO IOS specification.

Access channel packets can be routed in the first and second radio nodes based on a session identifier associated with the first and second access terminals. The session identifier associated with the first access terminal can be based on an access terminal identifier (ATI) based on a 1xEV-DO standard. The session identifier
20 associated with the second access terminal can be based on a temporary mobile subscriber identity (TMSI) based on a cdma2000 standard. The radio network controllers can include a PDSN function. A new session can be established for the first mobile access terminal on the second radio network controller or an inter-RNC dormant handoff from the first radio network controller to the second radio network controller
25 can be performed when the first access terminal is in dormant state after moving from a coverage area of the first radio node to a coverage area of the second radio node. The first radio network controller and the first radio node can be co-located. The second radio network controller and the second radio node can be co-located.

In one aspect, the invention features a method for exchanging digital
30 information using mobile access terminals in a wireless network. The method includes transmitting packets over a first traffic channel, established between a first mobile access terminal and a first radio network controller, via a first radio node without

passing through a second radio network controller, and transmitting packets over a second traffic channel, established between a second mobile access terminal and the second radio network controller, via a second radio node without passing through the first radio network controller. The method also maintains the first traffic channel
5 between the first access terminal and the first radio network controller when packets are received from or transmitted to the first access terminal via the second radio node. The method also includes transmitting packets received from the first access terminal to an external network (e.g., the Internet).

Implementations may include one or more of the following features. The
10 method may also include performing an active handoff of the first traffic channel from the first radio network controller to the second radio network controller upon detecting that the first access terminal is near a third radio node. The method may also include closing the first traffic channel by the first radio network controller upon detecting that the first access terminal is near a third radio node.

15 The method may also include establishing a first session for a third mobile access terminal on the first radio network controller via the first radio node, keeping track of the approximate location of the third access terminal based on location update messages sent by the access terminal, receiving packets at the first radio network controller for the third access terminal while it is in a dormant state, sending the second
20 radio controller a message requesting that it page the third access terminal, and paging the access terminal from the second radio network controller via a third radio node. The method may also include receiving at the third radio node a traffic channel request message from the third mobile access terminal over an access channel, forwarding the traffic channel request message from the third radio node to the second radio network
25 controller, performing a handoff of the first session from the first radio network controller to the second radio network controller, and, after completing the handoff, establishing a traffic channel between the second radio network controller and the third access terminal.

A radio network controller (e.g., the first radio network controller) may include
30 multiple server cards, and the method may also include establishing the first traffic channel between the radio network controller and the first mobile access terminal on one of a server card selected from the plurality of server cards, sending an address of

the selected server card from the radio network controller to the first radio node, sending reverse link traffic channel packets from the first radio node to the address of the selected server card, sending the address of the selected server card from the first radio network controller to the second radio node, and sending reverse link traffic
5 channel packets from the second radio node to the address of the selected server card.

The first radio node may be configured to operate on a first frequency channel while the second radio node may be co-located with the first radio node and configured to operate on a second frequency channel. The method may also include establishing a first session for the first mobile access terminal on the first radio network controller via
10 the first radio node operating on the first frequency channel, establishing a second session for the second mobile access terminal on the second radio network controller via the second radio node operating on the second frequency channel, and establishing a third session for a third mobile access terminal on the first radio network controller via the second radio node operating on the second frequency channel.

The method may also include establishing a primary association between the
15 first radio node operating on the first frequency channel and the first radio network controller, establishing a secondary association between the second radio node operating on the second frequency channel and the first radio network controller, establishing a first session for the first mobile access terminal on the first radio network
20 controller via the first radio node operating on the first frequency channel, and transferring the first session to the second radio network controller, whenever an access terminal starts monitoring the second radio node in a dormant mode.

The method may also include receiving a connection request from the first mobile access terminal in the first radio network controller via the first radio node
25 operating on the first frequency channel, and establishing a third traffic channel for the first access terminal on the first radio network controller, wherein the traffic channel flowing via the second radio node operating on the second frequency channel.

The method may also include establishing the third traffic channel flowing via the second radio node operating on the second frequency channel based on actual load
30 information received by the first radio network controller from the first and second radio node. The method may also include establishing a third traffic channel flowing via the second radio node operating on the second frequency channel based on an

estimate of load on the first and second radio nodes.

The first and second radio network controllers may each include a plurality of server cards, and the method may further include handling a plurality of traffic channels on a first server card in the first radio network controller, detecting an overload

5 condition in the first server card, selecting a traffic channel served on the first server card for transfer to another server card, and transferring the selected traffic channel to another server card in one of the radio network controllers without dropping the traffic channel.

Selection of the traffic channel for transfer may be based at least partially on the
10 amount of processing resources used by the selected traffic channel and/or on a quality-of-service requirement of the traffic on the plurality of traffic channels. Determination of a target server card to transfer the selected traffic channel to based at least partially on load and availability of other cards in the radio network controllers. A target server card may be located in the same or a different radio network controller as the first
15 server card.

Information about the load of server cards in one or more radio network controllers may be provided by a centralized load tracker that is located within a radio network controller or is external to all of the radio network controllers. A centralized load tracker may be configured to trigger a transfer of a traffic channel from the first
20 server card to the target server card. In some implementations, load information may be obtained by a server card directly from other server cards.

The method may also include using an Internet Protocol network to exchange data packets between the first radio network controllers and the first and second radio nodes.

25 In another aspect, the invention features a radio access network for wirelessly communicating with mobile access terminals that includes a plurality of radio nodes interconnected with a plurality of radio network controllers using a network (e.g., an IP network), wherein each said radio node can address each said radio network controller and each said radio network controller can address each said radio node, and an
30 interface for exchanging packets between the radio access network and an external network.

Implementations may include one or more of the following features. The radio

nodes and radio network controllers may be associated with a common sub-network. Each radio network controller may be configured to maintain a traffic channel regardless of which radio node the traffic channel is flowing. The radio network controllers may maintain a traffic channel by routing packets to and from the access
5 terminal via any one of the radio nodes. Each radio node may be associated with a primary radio network controller selected from the plurality of radio network controllers and each of the radio network controllers may be enabled to send a paging message to an access terminal via any of the radio nodes.

Each radio network controller may include a multiple server cards, each
10 connected to the network and addressable by each of the radio nodes. The radio network controller may be configured to establish traffic channels with access terminals on each of the plurality of server cards. The radio network controller may also be configured to provide an address of a server card on which a traffic channel is established to one or more radio nodes. Each radio network controller may also be
15 configured to detect an overload condition in one of its plurality of server cards and select one or more traffic channels handled by an overloaded card for transfer to another card

The plurality of radio nodes may include a two co-located radio nodes, with a first radio node configured to operate on a first frequency channel, and a second radio
20 node configured to operate on a second frequency channel.

The interface may be a packet data switching node or may be part of one of the radio network controllers.

Implementations can realize one or more of the following advantages. By establishing primary or secondary associations between all radio nodes and all radio
25 node controllers, a radio node controller has the information necessary to find and communicate with any radio node on the IP backhaul as an access terminal moves from the coverage area of one radio node to the coverage area of another radio node. This allows the radio network to perform normal soft handoff even when the access terminal moves into the coverage area of a radio node, whose primary RNC is different from the
30 one currently serving the access terminal. This avoids more traditional inter-RNC handoff procedures, which are known to be error-prone and introduce additional latencies. One implementation of the invention provides all of the above advantages.

DESCRIPTION OF DRAWINGS

Figures 1 through 4 show networks.

DETAILED DESCRIPTION

Existing 3G wireless network architectures, including the ones discussed above
5 for IS-856, assume a fixed relationship between RN's and RNC's. In other words, all
traffic flowing from or to an RN goes through the same RNC. This requires complex
hierarchical structures to handle dormant handoffs between RNC's and requires
frequent and delay-prone inter-RNC (soft) handoffs. Fixed associations between RN's
and RNC's are needed in circuit-switched voice applications when point-to-point
10 dedicated leased lines are used for backhaul connectivity between RN's and RNC's as
illustrated in Figures 1 and 2.

In the examples that follow, an RNC represents a radio network controller
directly attached to an IP network. Direct attachment means the RNC can send/receive
IP packets to/from other nodes in the IP network without any intermediate processing
15 above the IP layer. An RNC may be implemented on a chassis-based hardware
platform, which may consist of multiple server cards. In this case, the entire chassis
may be viewed as an RNC directly connected to the IP network, and the individual
server cards need an intermediate node (or an I/O server card) operating above the IP
layer to send/receive IP packets to/from other nodes in the IP network. In an alternative
20 implementation, individual server cards in a chassis-based hardware platform may be
directly attached to the IP network. In this case the server cards may have individual IP
addresses and they may each then be viewed as an RNC. An RNC may also be
implemented on standalone server hardware, such as a conventional compute server or
a blade server. In this case, the server may be viewed as the RNC and has an IP
25 address visible to other RN's and RNC's in the radio access network.

IP-Based Radio Access Network Architecture with Inter-RNC Signaling

First, as shown in Figure 3, consider the case where a set of RNC's 60 are co-
located in a data center and are connected together via a high-speed Local Area
Network (LAN) 62 such as a Gigabit Ethernet LAN. In this case, RNC's connect to the
30 network via LAN interfaces and a router 64 provides the connectivity to the external
network. Such a configuration can be referred to as an RNC cluster (or pool). (The
description below describes how the same concept can be extended to RNC's

connected over a more general IP network, such as a metropolitan-area network.) In the past, when the main traffic type carried through a wireless network was circuit-switched voice, such clustering using an Ethernet LAN was not feasible. RN's may connect to the router in the data center using dedicated leased lines 66. It is assumed
5 that RN's and RNC's are all IP addressable. In other words, any RN served by the cluster can communicate directly at the IP level with any of the other RNC's in the cluster.

In an RNC cluster such as the one described above, it is important to avoid any handoff boundaries between individual RNC's so that the entire cluster can behave as if
10 that cluster was one big RNC. This would eliminate unnecessary handoffs due to mobility, thereby greatly improving scalability and reliability.

To accomplish this, suppose in one example, an IS-856 subnet 70 is defined to be the entire footprint of the RNC cluster, not the footprint of just one RNC. In other words, all the RN's served by the cluster now belong to the same subnet. To simplify
15 the system operation, each RN in the subnet is associated (e.g., primary association) with one RNC in the cluster. This association is established when an RN is first powered. The detailed meaning of this association will be explained later.

Access Channel Packet Routing

Each sector in an RN can transmit to an AT over the forward traffic or control
20 channels 72. Similarly, each sector in an RN can receive from an AT over the reverse traffic or access channels 74. The access channel and the reverse traffic channels are separated by code division multiplexing using a Long Code Mask, whereas the control channel and the forward traffic channels are separated by time-division multiplexing using a preamble. The preamble identifies a forward link physical layer packet as a
25 control channel packet or as a traffic channel packet associated with a certain MAC Index. A MAC Index, an integer between 0 and 63, is unique within a sector and is assigned by the RN and RNC at the time of traffic channel establishment. Similarly, the Long Code Mask identifies a reverse link physical layer packet as an access channel packet or a specific traffic channel packet. The Long Code Mask is based on the AT's
30 UATI for the traffic channel, and is based on the SectorID of the serving sector for the access channel. The sending AT of an access channel packet and the recipient AT of a control channel packet are indicated in the ATI field of a MAC Layer header.

Whenever an RN receives a MAC Layer packet on one of its access channels, the RN forwards the packet, without even looking at its content, to its primary (or Default) RNC in the cluster with whom that RN is associated. As such, when a packet carrying a UATI_Request message is received from an AT, the packet is forwarded by the receiving RN to the Primary RNC. The RN encapsulates the MAC Layer packet in an IP packet (possibly multiplexed with MAC Layer packets of other AT's) with a destination IP address equal to an IP address of the serving RNC. The IP packet is carried over the backhaul network to an aggregation router at the data center and the router forwards the packet to the serving RNC over the Ethernet LAN.

All access channel packets include an address field that identifies the sending AT. When the sending AT has already been assigned a UATI by an RNC, the address field contains that UATI. When the sending AT does not yet have a UATI, the address field contains a Random Access Terminal Identifier (RATI), which is randomly selected by the AT. The first two bits of the address field indicate whether the address is a UATI or a RATI.

When the (Ethernet) I/O subsystem of an RNC receives a UATI_Request message from an AT with an address field that contains a RATI or an unrecognized UATI, the RNC assumes the role of the serving RNC to handle the session. If the RNC is implemented on a chassis-based hardware platform, it assigns the session to one of its server cards. The AT is then assigned a UATI within some predetermined range. This range, which identifies the serving RNC to all other RNC's in the cluster, is known by all the RNC's in the cluster, but is not known by the AT. If the RNC is implemented on a chassis-based hardware platform, the range of the UATI's that belong to a certain RNC may further be subdivided to identify the server card within the serving RNC that is handling the session. The serving RNC also establishes an A10 connection with the PDSN in order to facilitate the data transfer between the AT and the PDSN. The A10 connection terminates on the server card handling the session.

While dormant, the AT sends RouteUpdate messages, as needed, to provide information about its current location. This mobility information is maintained at a Mobility Manager in the serving RNC. Since a subnet covers the entire footprint of the RNC cluster, when the AT crosses the boundary between two RNC's in the same

cluster, the AT does not detect a subnet change and therefore does not initiate a dormant handoff. But when the AT sends an access channel message to an RN that is associated with a different RNC (broker RNC) in the cluster, the packet(s) carrying that message are sent by the RN to the broker RNC. The I/O subsystem in the broker RNC
5 examines the address field of all arriving access channel packets and reads the UATI. From the UATI, the I/O subsystem determines by table look-up the identity of the serving RNC and forwards the access channel packet to that RNC over the high-speed LAN. When the UATI on the access channel packet is served by the receiving RNC, it processes the packet locally. If the receiving RNC is implemented on a chassis-based
10 hardware platform, its I/O subsystem first determines the server module (card) that is handling the UATI (session) and forwards the packet to that card using an internal bus of the serving RNC.

Page Routing

If packet data is received from the PDSN for a dormant AT, the packets are
15 forwarded over the A10 interface to a specific server card on the serving RNC. That server card then looks up the location information for that AT. The serving RNC then sends a paging message via a set of RN's that are determined based on the last Route Update message received from the AT. The paging message is sent via the control channel of one or more sectors that belong to the RNC cluster. The RN's transmitting
20 the paging message may not be associated with the serving RNC (i.e., they may have a different Primary RNC), but they need to be associated with one of the RNC's in the cluster.

Connection (Traffic Channel) Establishment

When a serving RNC receives a ConnectionRequest message from the AT,
25 either directly or via a broker RNC, it examines the pilot strengths reported by the AT in the RouteUpdate message accompanying the ConnectionRequest message. To simplify system operation, it is assumed that each RN's radio resources are managed by a Radio Resource Control function in the RNC with whom the RN is associated. Furthermore, an RN can exchange signaling only with the RNC with whom it is
30 associated. Therefore, when the serving RNC wants to establish a traffic channel that involves RN's that are associated with other RNC's, the serving RNC first communicates directly with the Radio Resource Control function on those RNC's to

check for resource availability. Such communication occurs over the high-speed LAN. (The serving RNC may use a look-up table to determine the RN's primary RNC.)

When sufficient radio resources are available, the serving RNC establishes the necessary traffic channel communication links with the RN's via the RNC's with whom
5 these RN's are associated and sends a TrafficChannelAssignment message to the AT to initiate the traffic channel set up. Once a traffic channel has been established, packets flow directly between the RN's and the serving RNC without any involvement of any broker RNC. Such direct routing eliminates the delays typically found in soft handoff procedures that involve triangular routing through another RNC.

10 When a new traffic channel involves an RN that is outside the footprint of the RNC cluster (different subnet), a similar procedure is implemented. In this case, the serving RNC communicates with the RNC's outside the cluster over an IP network (a metropolitan-area network) to obtain radio resources. If the radio resources are available, the serving RNC establishes a communication link with that RN by
15 exchanging signaling via the RNC's with whom these RN's are associated.

This method allows the serving RNC to maintain the traffic channel, even when the AT moves into the coverage area of an RN that is associated with an RNC different from the serving RNC.

Improved IP-based Radio Network Architecture Without Inter-RNC

20 Signaling

The scheme described so far can be improved in couple of areas. First, the triangular routing of access channel packets via a broker RNC can be eliminated by moving that routing function to the RN's. This will reduce delays in handling access channel packets, for example during traffic channel set-up, at the expense of some
25 increase in processing power at the RN. Furthermore, one can eliminate all signaling between RNC's and instead allow RNC's directly exchange signaling with all RN's in the radio access network. This helps create a simpler network architecture, which is easier to deploy and maintain.

The Radio Resource Control function can also be moved from the RNC's to the
30 RN's, thereby further reducing delays in traffic channel set-up procedures.

The IP-based radio access network and its enhanced version described here both allow exploitation of the flexibilities of IP and metropolitan Ethernet networks and

result in a more distributed system where an AT may remain attached to its serving RNC regardless of its position, except when the distance between the AT and the serving RNC becomes excessive. To provide these capabilities each RN may associate with multiple RNC's, possibly with all the RNC's in the IP-based radio access network,
5 but here it is not necessary to have one of the RNC's behave as a primary RNC.

Avoiding Triangular Routing of Access Channel Packets

When powered on for the first time, an AT registers with the IS-856 network as follows: The AT acquires an IS-856 pilot being broadcast by one of the nearby sectors and synchronizes with the system. To initiate the session establishment, the AT sends a
10 UATI_Request. As before, the AT uses a Random ATI (RATI) in the MAC Layer header to send this request.

The RN examines the address field of the access channel packet and recognizes that the originator of the message does not have an assigned UATI and forwards the packet to its Primary RNC with whom the RN is associated. To examine the address
15 field, the RN first extracts the MAC Layer capsule fragments from the received MAC Layer packets, and forms the MAC Layer capsule. The RN then reads the address field in the MAC Layer header.

After receiving the UATI_Request, the Primary RNC assumes the role of the serving RNC and assigns a UATI to the AT. The primary RNC then proceeds with the
20 rest of session establishment, in particular the security key exchange and the protocol configurations. (An improved version of this procedure to increase availability and to provide better load balancing is described in more detail below.) The RNC also implements the PPP/CHAP procedure to authenticate the AT based on its Network Access Identifier (NAI). There is a one-to-one mapping between the NAI and the
25 terminal's actual IMSI (International Mobile Subscriber Identity). This mapping is maintained in an AAA (Radius) server (not shown). The AAA server passes the AT's IMSI value to the serving RNC.

The pocket control function (PCF) function in the serving RNC uses this IMSI value to select a PDSN as described in the IS-2001 standard and establishes an A10
30 connection to that PDSN. In the A11 Registration message, the PCF function provides the IMSI value of the AT along with its own SID/NID/PZID identifier to the PDSN. The AT and the PDSN then set up a PPP link, perform Simple IP or Mobile IP set-up

and execute user-level authentication.

Each RN keeps a routing table for the mapping between the UATI and the serving RNC. This routing table may be provided to the RN by a network management system. As in the previous system, each RNC owns the UATI values that fall within a certain range. Whenever the RN receives an Access Channel packet, the RN
5 determines from the UATI value in the MAC Layer Header the identity of the serving RNC, and routes the packet to that RNC by placing an IP address of the serving RNC in the destination address field of the IP header. Thus access channel packets are delivered by any RN directly to the serving RNC. The range of UATI values that each
10 RNC owns can be communicated by the RNC to the RN's directly, thereby eliminating the need for explicitly configuring the RN's from the management system with UATI ranges.

The above method requires that the RN maintains a table for mapping the UATI to the IP address of the RNC. Alternatively, the UATI space can be divided among all
15 the RNC's, and each RNC be assigned a unique subspace, and an algorithmic relationship can be established between the IP address of the RNC and the UATI subspace. The RN can then determine the IP address of the RNC from the AT's UATI algorithmically, without using any tables.

It is also possible to have hybrid schemes where access channel packet routing
20 is handled in both a central element and also in a distributed fashion in the RN's. Here, a central UATI server, possibly located in one the RNC's, may be responsible for assigning UATI's and the associated new sessions to serving RNC's. When first establishing a new session, an RNC can request a new UATI from the central UATI server. The serving RNC can also register with a centralized AC router who establishes
25 a binding between the UATI and the serving RNC. When an RN is serving an AT for the first time, it can forward the packet to the UATI router, who can then forward the packet to the UATI router. The serving RNC can then perform a binding update with the serving RN, so that in all subsequent transactions access channel packets can be sent directly to the serving RNC, avoiding triangular routing.

30 **Avoiding Handoffs Between RNC's**

As before, mobility management for a given AT is handled entirely by the serving RNC. The AT is configured to provide distance-based location update in

dormant mode. In other words, whenever the serving sector is more than a certain distance away from the sector where it last sent a RouteUpdate message, the AT sends a new RouteUpdate message to the serving sector over the Access Channel. The RouteUpdate message is forwarded by the RN to the serving RNC which then keeps
5 track of the location of the AT.

When the serving RNC wants to page an AT, the serving RNC first determines the RN or RN's from which it wants to send the page, depending on the time and position indicated in the most recent RouteUpdate message received from the AT. It is assumed here that the serving RNC knows the IP addresses of all the RN's in the radio
10 access network. The RNC can obtain this information from a network management system or directly from the RN's when they associate with the RNC during initial power-up. The serving RNC sends the paging message to the appropriate set of RN's directly. These RN's then page the AT over their respective control channels.

All sectors in an IS-856 network broadcast in their overhead channel their
15 SectorID and Subnet Mask. For a relatively small network, one can set the Subnet Mask to zero, thereby implying that the entire network is one big subnet. In this scenario, the AT never detects a subnet change. Therefore, the AT remains attached to the original serving RNC and does not trigger a dormant inter-RNC handoff. The A10 connection to the PDSN also remains fixed regardless of the position of the AT.

20 If the radio access network covers a geographically large area, it may be prudent to force a dormant inter-RNC handoff, when the AT moves too far away from the serving RNC. This can be triggered by the serving RNC, for example, upon receiving a RouteUpdate message from the AT. Alternatively, the Subnet Mask can be chosen greater than zero to introduce a subnet boundary between too far-away RNC's. Then,
25 when the AT crosses the subnet boundary, a dormant handoff occurs and the A10 connection is relocated. Further, the AT is assigned a new UATI and session parameters are transferred from the old serving RNC to the new serving RNC.

Faster Traffic Channel Setup Using Distributed Radio Resource Control

What follows next is a description of how moving the Radio Resource Control
30 from the RNC's to the RN's and establishing a direct signaling link between all the RNC's and the RN's in the radio access network reduces the set-up time for traffic channels that (in the previous scheme) involved multiple RNC's. Whenever the AT

sends a ConnectionRequest message over the access channel along with a RouteUpdate message to initiate a new traffic channel, the message is immediately forwarded from the receiving RN to the serving RNC. The serving RNC examines the RouteUpdate message to determine a likely set of sectors that may be included in the Active Set. The
5 serving RNC then corresponds directly with the RN's where these sectors reside, to request traffic channel and backhaul resources. The RN's either decline, or accept and allocate the needed radio resources. If resources are available from a sufficient set of RN's, the serving RNC accepts the traffic channel request, and sends a TrafficChannel assignment message over the Control Channel to the AT. The AT then starts
10 transmitting on the Reverse Traffic Channel (RTC). Once the RN acquires the RTC, an RTCAck message is sent to the AT to indicate the acquisition of the RTC signal. The AT then responds with a TrafficChannelComplete message to indicate the completion of the Traffic channel set-up.

In this procedure each RN controls its own radio resources, both with respect to
15 hardware resources available on the RN, as well as the management of interference across its sectors. As a result, the admission control function is split between the RN and the serving RNC. RN's provide local admission control for the sectors they control while the serving RNC provides a global admission control. Similarly, when a sector in a given traffic channel is inactive for some period of time, the sector can initiate the
20 procedure for closing the traffic channel by sending a request to the serving RNC to close the traffic channel. The serving RNC then makes a global decision on whether to remove that sector from the traffic channel, close the entire traffic channel or do nothing.

Once a traffic channel has been setup between the AT and a serving RNC, it
25 remains anchored to the serving RNC, even when the AT moves into the coverage area of other RN's in the IP-based radio access network.

Packet Routing Between RN and RNC - In More Detail

When a sector in the RN receives a MAC Layer packet on a reverse traffic channel, the sector forwards the packet to an I/O card after adding a Stream Identifier
30 that includes a Connection Identifier. The I/O card uses the Connection Identifier value to look up the IP address of the serving RNC. The I/O card then encapsulates the MAC

Layer packet together with its Stream Identifier in an IP packet whose destination address is set to the IP Address of the serving RNC. If the serving RNC is implemented on a chassis-based hardware platform, the I/O module in the serving RNC, upon receiving the packet, reads the UATI value to determine the server module
5 that handles this session. The I/O card then passes the packet along with the Stream Identifier to that server module for further processing.

When a sector in the RN receives a MAC Layer packet on the access channel, the sector first reads the UATI in the ATI field of the MAC Layer Header and then forwards the packet to an I/O card after adding a Stream Identifier that includes the
10 UATI of the sending AT along with the serving sector's SectorID. The I/O card in the RN again uses the UATI value to look up the IP address of the serving RNC. The I/O card encapsulates the MAC Layer packet together with its Stream Identifier in an IP packet whose destination address is set to the IP Address of the serving RNC. If the RNC is implemented on a chassis-based hardware platform, the I/O module in the
15 serving RNC, upon receiving the packet, reads the UATI value to determine the server module that serves this session. The I/O card then passes the MAC Layer packet along with the Stream Identifier to that server module for further processing.

When the serving RNC has a MAC Layer packet ready for transmission on a forward traffic channel, it first looks up the IP address of the RN to which to send the
20 packet. It then encapsulates the MAC Layer packet together with a Stream Identifier in an IP packet whose destination address is set to the IP Address of the RN. The RN, upon receiving the packet, reads the SectorID value in the Stream Identifier to determine the sector that will transmit the packet. The RN then passes the MAC Layer packet along with the Stream Identifier to the appropriate modem card, which
25 schedules the MAC Layer packet for transmission on the Forward Link using the MAC Index as the preamble.

Similarly, on the forward link, when a serving RNC has a MAC Layer packet ready for transmission on the Control Channel of a particular sector, the serving RNC determines the IP address of the RN to which to send the packet. It then encapsulates
30 the MAC Layer packet together with its Stream Identifier in an IP packet whose destination address is set to the IP Address of the RN. The RN, upon receiving the packet, reads the SectorID value in the Stream Identifier to determine the sector that

will transmit the packet. The RN then passes the MAC Layer packet along with the SectorID and MAC Index to the appropriate modem card. The modem card schedules the packet for transmission on the control channel.

Failure Recovery & Load Balancing

5 The improved IP-based radio access network architecture described earlier can be further extended to increase the overall reliability of the wireless network.

Failure Recovery Without Session Preservation

First, consider an approach where each RN, upon power-up, first communicates with a primary RNC Resource Control Agent who may reside in one or more of the RNC's or may reside on a separate compute engine or server. The primary Resource Control Agent assigns each RN to a Primary RNC. The RN then routes all new session requests to that Primary RNC.

When an RNC becomes completely unreachable due to some failure, all AT's that are being served by that RNC will ultimately recognize that their IS-856 sessions have been lost. Each of these AT's will initiate a new session by sending a UATI_Request over the Access Channel. Every RN who receives one of these requests will route them to its primary RNC. If at any time, the RN cannot reach its primary RNC, the RN will immediately request a new primary RNC from the primary RNC Resource Control Agent. If the primary RNC Resource Control Agent is also not reachable, the RN will send a similar request to a secondary RNC Resource Control Agent. Once the UATI_Request is received by the Primary RNC, the Primary RNC will immediately establish a new IS-856 session with the AT and will further initiate the procedure to set up a new A10 connection with a PDSN.

Assignment of a new Primary RNC may also be initiated by the RNC Resource Control Agent. This can be accomplished by having the RNC Resource Control Agent continuously monitor the health of all the RNC's in the subnetwork. Upon detecting the failure of an RNC, the RNC Resource Control Agent immediately communicates with all affected RN's and assigns them to new Primary RNC's. In assigning RN's to Primary RNC's, the RNC Resource Control Agent may perform load balancing to ensure that user sessions are evenly distributed across all available RNC's.

Load Balancing Session Assignment

The above method can be further enhanced by making the RNC Resource

Control Agent ultimately responsible for assigning user sessions to RNC's. In this case, when a primary (or Default) RNC or possibly the RN itself receives a new UATI_Request, the Primary RNC (or the RN) asks the RNC Resource Control Agent to assign the session to an RNC. The RNC Resource Control Agent assigns the session to
5 an RNC based on resource availability, loading and the distance between the RNC and the RN presently serving the AT. This approach provides better load balancing among RNC's, allowing user sessions to be distributed across RNC's more dynamically, while also taking into account the current position of the AT. In case of an RNC failure, all new session requests will arrive at the RNC Resource Control Agent who will then
10 assign these sessions to new RNC's, again based on loading and other considerations.

The RNC Resource Control Agent may also be used to trigger dormant handoffs for load balancing or other purposes. In Phase 1 IS-856 networks, a dormant inter-RNC handoff is always triggered by the AT upon detection of a subnet change. As described above, lack of an immediate dormant handoff may result in lost paging data.

15 In the improved IS-856 networks shown in Figures 3 and 4, a dormant handoff can be initiated by the network based on the location of the AT. Upon receipt of a RouteUpdate, when a serving RNC determines that a transfer of a user session to another RNC is desired (for load balancing or other reasons), the serving RNC sends a Dormant Handoff request to the RNC Resource Control Agent who assigns the session
20 to a new RNC. The new serving RNC then assigns a new UATI and performs a session transfer from the previous serving RNC.

In a more distributed implementation of the RNC Resource Control Agent concept, RNC's can constantly communicate with the RN's and other RNC's to provide routing information (including their loading) to all the RN's, thereby allowing
25 the RN's to route incoming session requests to the correct RNC without going through a RNC Resource Control Agent. For example, each RN may have a preferred RNC list and whenever it needs to assign a new session to an RNC, it selects one of the RNC's in this list according to some algorithm (pseudo-random selection, round-robin, etc.) If an RNC in the preferred list were to become unavailable, the RN would detect this
30 (KeepAlive messaging can be used between RN's and RNC's to help the RN detect an RNC failure), and remove that RNC from its preferred list. A drawback of this approach is that some backhaul signaling traffic would be created as a result of

exchanging such dynamic loading information.

Failure Recovery with Session Preservation

In some networks, it may be necessary to recover user session information in case of an RNC failure. This would eliminate the air link congestion that hundreds of
5 new session requests could create shortly after an RNC failure. In order to preserve sessions in case of failure of an RNC, a copy of such information (for all sessions in the subnetwork) can be stored in the RNC Resource Control Agent.

When an RNC fails and the AT initiates a new session, its new session request will reach the RNC Resource Control Agent. The RNC Resource Control Agent then
10 not only assigns a new serving RNC to each session, but also provides the session information thereby avoiding lengthy session establishment procedures. Once a new UATI is successfully assigned to the AT, communication with the network may resume. The RNC Resource Control Agent further provides information related to the A10 interface, in order to allow the RNC establish an A10 session with the same
15 PDSN, thereby avoiding the setting up of new PPP and Mobile/Simple IP sessions.

In a chassis-based RNC, the RNC Resource Control Agent may run on a specific redundant card, with a hot standby. The RNC Resource Control Agent is then responsible for storing session information. In case a server module fails, the session is internally reallocated to another server module. In principle, the operation of this
20 system is the same as the one operating across the network. Moreover, in this case, it is not necessary to reestablish the A10 session to the PDSN, since the external IP address of the PCF seen by the PDSN can be maintained.

Integrated RNC & PDSN

Another benefit of the IP-based radio access network architecture described
25 above is the ability to combine the RNC and PDSN functions in a single network element. In hierarchical 3G packet data networks, a PDSN represents the highest point in the hierarchy, and therefore can support multiple RNC's. A new generation of PDSN's are expected to support hundreds of thousands of users, and several RNC's.

In existing radio access networks with dedicated point-to-point links between
30 RN's and RNC's, migrating the PDSN function to the RNC would be undesirable, because this would reduce the number of sessions that could be supported, resulting in frequent costly handoffs between PDSN's that involve new PPP and Simple/Mobile IP

registrations.

In the IP-based radio access network architecture described here handoffs between RNC's occur much less frequently therefore allowing the integration of the PDSN function into the RNC. Since an active call can always be served by the same
5 RNC, inter-PDSN handoffs are normally not required during an active call. Such an approach also simplifies the networking between the RNC and the PDSN, and further increases scalability and reliability.

In an RNC with an integrated PDSN, the PDSN functionality includes PPP termination, Simple IP and/or Mobile IP foreign agent and the AAA client functions.
10 As long as the AT remains within a subnet (say an RNC cluster), no inter-PDSN handoffs would be required.

If an integrated RNC/PDSN fails, all sessions supporting an AT (including the air interface, PPP and Simple/MobileIP sessions) are transferred to another RNC/PDSN thereby avoiding any new session establishment between the AT and the wireless
15 network.

It is also possible to integrate the RNC/PDSN with the RN. In this case, the RNC/PDSN functionality may be co-located with the RN in the same site, or even in the same enclosure.

It should be understood that the methods described in this disclosure are equally
20 applicable to networks where the RNC and the PDSN or the RN, RNC and the PDSN are integrated or co-located.

IP-based Radio Access Network Architecture with Primary/Secondary RNC Association

Each RN has a primary (e.g., default) RNC with which the RN associates (e.g.,
25 establishes a primary association), as described above. Using the backhaul network 80, each RN may also associate with one or more other RNC's in the IP RAN, and these RNC's are referred to as secondary RNC's for that RN (e.g., establishes a secondary association). To associate with an RNC, the RN provides sufficient information about itself to the secondary RNC's to allow these RNC's to communicate with AT's via that
30 RN. Further, to support signaling exchanges between the RN and secondary RNC's, signaling connections are established between them. The RNC also distinguishes between primary RN's and secondary RN's. Primary RN's are those that have a

primary association with that RNC. The establishing of secondary associations can be performed, for example, when an RN powers up.

When an RN receives an access channel packet from an AT, the RN forwards the packet blindly to its primary RNC. A dormant AT in the coverage area of an RN is served by the primary RNC of that RN. In a basic implementation, all RN's that share the same primary RNC belong to the same 1xEV-DO subnet. When the AT crosses a subnet boundary, the AT sends a UATI-Request to its serving RN, who then forwards the request to its primary RNC. Since this UATI is not served by that RNC, that RNC initiates the normal procedures for dormant inter-RNC handoff with the old serving RNC according to the A13 interface defined in 1xEV-DO IOS. The new RNC also performs an A10 handoff. In summary, this method does preserve the dormant handoff boundaries between RNC's.

It is also possible to allow RN's that belong to different 1xEV-DO subnets to share the same primary RNC. In this case, when the AT sends a UATI-Request upon crossing a subnet boundary, the message gets forwarded to the same primary RNC that is currently serving the AT (because the same RNC is the primary RNC for RN's in both subnets). The primary RNC, recognizing that it is already serving this UATI, can then proceed with a UATI assignment without an A13 procedure or an A10 handoff.

When the primary RNC receives an incoming data over an A10 connection for a dormant AT, the primary RNC proceeds with the paging procedure as usual. The paging message is again sent via all RN's that have this RNC as the primary RNC, or a subset thereof.

When a serving RNC receives a traffic channel request message from an AT requesting the set up of a traffic channel, the serving RNC first examines the Route Update message to determine the pilots that are needed and finds the RN's where these pilots are located. The serving RNC then contacts these RNs to set up the traffic channel as usual. The primary/secondary association allows the RNC to map the PN offsets of the requested pilots to RN's IP addresses and use the pre-established signaling connection with these RN's to set up all handoff legs. Once an active traffic channel is established, the serving RNC remains anchored as the RNC for that traffic channel. The serving RNC adds and removes RNs as the AT moves through coverage areas using the secondary associations of the RNs.

For example, as an AT first communicates with a first RN, that RN, as described above, forwards a request to its primary RNC (i.e., the RNC with which the first RN has a primary association). Once a traffic channel is established, the primary RNC remains the serving RNC for that entire traffic channel as the AT travels from the coverage area of the first RN to the coverage area of a second RN, even though that second RN does not have a primary association with the serving RNC. This allows the user activity (e.g., phone call, data transfer) to continue uninterrupted. While the AT is moving from the coverage area of the first RN to the coverage area of the second RN, the AT communicates the pilot strength information for the second RN to the serving RNC by sending a RouteUpdate message. As in normal soft handoff between RN's that have the same primary RNC, this communication occurs while the AT is still using the first RN. Upon reception of the RouteUpdate message, the serving RNC can communicate with the second RN because of the established secondary association, which, as described above, includes information such as the PN offset of the pilot signal and the IP address of the RN, to allow the serving RNC to contact the second RN and establish a communication channel. One appealing aspect of the handoff procedure above is that it works just like normal soft handoff. The only difference is the secondary association between RN's and RNC's, which allow the RNC to set up a traffic channel with RN's with whom it does not have a primary association.

Since RN's always forward their access channel packets to their primary RNC, there is no longer a need for the RN to perform access channel packet routing based on UATI. This greatly simplifies the implementation of the IP-based radio access network.

Chassis-Based Systems

The concepts described above can also be used in chassis-based systems by, for example, treating each server card as an IP addressable RNC or each modem card as an IP addressable RN. The logical operation of the system is unchanged.

However, in large networks with many server cards, treating each server card as an independent RNC may result in too many RNC's, which in turn can create many signaling connections and RN-to-RNC associations. One way to hide this complexity is to treat the entire chassis as an IP addressable RNC and handle intra-chassis communication using internal protocols. Again the concepts described above can be

used in the same fashion, and the logical operation of the RNC and RN are unchanged when viewed from the outside. The following description describes in more detail how these concepts impact the internal operation of a chassis-based RNC.

In some of the examples above, when an RN associates with its primary or
5 default RNC, a signaling connection is established between the RN and RNC. In a chassis-based system, one of the server cards may assume the responsibility for performing the association and terminating the signaling connection. This means, when a new RN wants to associate with an RNC, the RN is internally assigned (homed on) to one of the server cards. Subsequently, all signaling between the RNC and the
10 RN is performed via that server card. In the IP-based radio access network architecture with inter-RNC signaling, a chassis-based RNC can discover other RNC's, including the PN offsets these RN's serve and the UATI space they control, using an inter-RNC topology manager, which may reside in a system controller (SC) card. System controller card communicates with the server cards to learn the necessary information
15 about the RN's to pass onto other RNC's.

In some implementations in which a chassis-based RNC is utilized, when an RN forwards a received AC packet to a chassis-based RNC, the packet is first intercepted by an I/O card, which inspects the UATI address in the AC packet, determines the server card serving that UATI and forwards the packet to that server card. When the
20 UATI is served locally, the I/O card determines the specific server card, which is serving the UATI and forwards the packet to that server card. To perform such forwarding, the I/O card maintains tables that map UATI addresses to RNC's and further maps local UATI addresses to server cards.

When a server card in a chassis-based RNC receives a request for a traffic
25 channel set up from an AT, the server card first determines from the RouteUpdate message the set of RN's whose pilots are needed for the active set by mapping the received PN offsets to RN IP addresses. The server card then exchanges signaling directly with those RN's that are homed on that server card to establish the needed traffic channel legs. For other RN's that are homed on the same RNC chassis, if there
30 are any, the server card first determines the server cards where those RN's are homed. The server card contacts these server cards sharing the same chassis, which in turn contact the RN's to establish a traffic channel leg. For other RN's that are homed on

other RNC chassis, the server card contacts these RNC's. The I/O card in these RNC's then route the request to the server cards where the RN's are homed on. These server cards then contact the RN's to establish the needed traffic channel legs.

Since direct IP communication is possible between all RN's and RNC's in the
5 IP RAN, all traffic channel legs operate directly between the RN's and the serving RNC. When the I/O card of a chassis-based RNC receives a traffic channel packet, the I/O card forwards the packet to the server card where the traffic channel is being handled based on a unique connection identifier.

In some implementations, each server card can be assigned an IP address and it
10 can supply this address to the RN's that are serving the traffic channel. The RN's can then send traffic channel packets directly to the serving RNC. The packets are then routed via the I/O card to the server card without requiring a higher-layer look-up based on a connection identifier.

The chassis-based RNC can maintain communication with an AT, even when
15 the AT is in the coverage area of an RN for which it is not the primary RNC. All user traffic flows directly between the RN's and the serving RNC. When a dormant AT travels across the entire IP RAN, an inter-RNC handoff is not required (although possible), since the serving RNC can page an AT via any RN in the IP RAN.

In the IP-based radio access network architecture with inter-RNC signaling,
20 there is discovery and communication between the RNC's via the System Controller cards, as described above. This architecture can also result in triangular routing of access channel packets when the primary RNC is not the serving RNC, and packets are routed via the I/O card of another chassis-based RNC. Also, the failure of an RNC server card results in the failure of all RN's who are homed on it. If the entire RNC
25 chassis fails, all primary RN's are lost.

In the improved IP-based radio access network architecture without inter-RNC signaling, the RN associates with multiple RNC's in the IP RAN and maintains a table to map a UATI address to the serving RNC. In this approach, RN's may be configured with the subnet information, while in the other methods with a primary RNC this
30 information may be supplied by the primary RNC. The RN also maintains a signaling connection with multiple RNC's in the IP RAN. All signaling now terminates in one of the server cards on the chassis-based RNC, which is assigned to home that RN. The

RN provides information about itself, including its IP address to the server card where it is homed. This server card in turn distributes the information to all other server cards in the chassis-based RNC.

In these examples, when the RN receives an access channel packet from an AT, the RN inspects the UATI address field, determines the IP address of the serving RNC from its table, and then forwards the packet to that RNC. The I/O card in the chassis-based RNC intercepts the packet, identifies the packet as an access channel packet, inspects the UATI address, and routes the packet to the server card handling this AT. In traffic channel set up, the server card directly contacts those RN's, which are homed on the server card to request the set up of a traffic channel leg. For other RN's, the server card first determines the other server card where the RN is homed, and contacts that other server card who in turn contacts the RN to set up the traffic channel leg. In this approach, RNC's can behave as autonomous entities with no inter-RNC communication, except to handle possible inter-RNC handoffs to optimize routing. However, there is inter-card communication to handle homing.

Even though triangular routing between RNC chassis have been eliminated there can be indirect routing inside the chassis, if all packets travel via the I/O card, which performs a higher-layer routing function. The failure of a server card results in the failure of the RN's that are homed on it. These RN's can no longer serve any of the sessions that reside on the RNC, even those that are still alive on server cards other than the failed server card. Such problems do not occur when each server card in the chassis behaves as an IP addressable RNC. In this case, packets are routed directly between RN and RNC server cards and failure of a server card does not impact RN's or other server cards.

In IP-based radio access network architecture with primary/secondary association, a set of chassis-based RNC's are connected to a set of RN's via an IP backhaul network and there are primary or secondary associations between all of the RNs and all of the RNCs.

Each RN has a primary chassis-based RNC with which the RN associates as described above. Each RN also associates with other chassis-based RNC's in the IP RAN, and these are referred to as secondary RNC's for that RN. The RN provides sufficient information about itself to the secondary RNC's to allow these RNC's to

communicate with AT's via that RN. Further, to support signaling exchanges between the RN and secondary RNC's, signaling connections are established between the RN and its secondary RNC's. In each case, the association between the RN and the chassis-based RNC is handled via one of the server cards in the RNC chassis.

5 Access channel packets that are blindly forwarded by the RN to its primary RNC are intercepted by the I/O card, which inspects the UATI field and forwards the packet to the server card currently handling that UATI. A dormant AT in the coverage area of an RN is served by a server card in the primary RNC of that RN. All RN's that share the same primary RNC belong to the same 1xEV-DO subnet. When the AT
10 crosses a subnet boundary, the AT sends a UATI-Request to its serving RN, who then forwards the request to its primary RNC. Since this UATI is not served by that RNC, the I/O card after inspecting the UATI field, forwards the request to any one of its server cards. A server card can then initiate the normal procedures for dormant inter-RNC handoff with the old serving RNC according to the A13 interface defined in
15 1xEV-DO IOS. The new server card also performs an A10 handoff.

It is also possible to allow RN's that belong to different 1xEV-DO subnets to share the same primary RNC. In this case, when the AT sends a UATI-Request upon crossing a subnet boundary, the message gets forwarded to the same primary RNC that is currently serving the AT (because the same RNC is the primary RNC for RN's in
20 different subnets). By inspecting the UATI address, the I/O card in the primary RNC can forward the packet to the server card that is currently serving that UATI, and that server card can then proceed with a UATI assignment without an A13 procedure or an A10 handoff.

When the primary RNC receives incoming data over an A10 connection for a
25 dormant AT, the packet is again intercepted by the I/O card, which, after inspection forwards the packet to the server card handling that A10 link. The server card then proceeds with the paging procedure as usual. The paging message is again sent via all RN's that have this RNC as the primary RNC, or a subset thereof.

When a server card receives a traffic channel request message from an
30 AT requesting the set up of a traffic channel, the server card examines the Route Update message to determine the pilots that are needed and finds the other server cards that are currently handling these pilots. The server card contacts these other server

cards, which in turn contact the RNs to set up the traffic channel as usual. The primary/secondary association allows the server card to map the PN offsets of the requested pilots to identifiers for the server cards where the RN's are homed and use the pre-established signaling connection through these homed RN's to set up all
5 handoff legs.

The description above can be further extended to increase reliability in case of a server card failure. In some examples, when a server card fails, all user sessions served by that card are also lost. This means the user may remain unreachable for a considerable amount of time. The user may also be unaware that he/she is unreachable.
10 To prevent this, when a user session is first established a copy of the session state information is stored on a separate card, for example, the system controller card. Session information includes various protocol configurations, mobility information, UATI, etc. When the session parameters change, the session information on the system controller card is updated.

15 The system controller card uses a heartbeat signaling mechanism to detect the failure of a server card and assigns the session to any one of the remaining server cards. The new server card may reassign the UATI to ensure that the new server card has connectivity with the AT. In this method, using just one extra server card in a load-balanced N+1 redundancy configuration will ensure that there is enough headroom to
20 reassign the sessions on the failed server card.

Even though some of the techniques described above employ the 1xEV-DO air interface standard, the techniques are equally applicable to other CDMA and non-CDMA air interface technologies. In this case, the link layer address (ATI) could be something different from ATI used in 1xEV-DO (e.g., TMSI in cdma2000), and other
25 RN zones, such as paging zones, PCF zones, etc, could be used instead of 1xEV-DO subnets.

Connection-Level Load Balancing and Overload Control in Chassis-Based RNC

In a chassis-based RNC, if loading on a server card exceeds processing and
30 memory capabilities of the card, the performance of all users being served by the card may be affected. To prevent such overload conditions, connection-level load balancing may be used.

For example, when processing and memory usage becomes too large, the server card may trigger an active session transfer for one or more connections to one or more other cards where resources may be available. In such load balancing or overload control schemes, it is desirable to have some central entity in the chassis, e.g., a system controller card, to keep track of the individual loading of each server card on the chassis. In some implementations, when a server card becomes overloaded, server card determines which connection to transfer to another card. Here the server card may use one of several criteria, including the QoS needs of individual flows in the connection, how much processing and memory resources the connection has recently been using, the QoS-class of the user, etc. Once the server card determines which connection(s) it wants to transfer to another server card, it contacts a central load tracker associated with the RNC and requests target server card(s) where these connections can be transferred.

When the central load tracker provides a list of available server cards, the overloaded card directly contacts one or more of the available server cards to initiate an active session transfer (or handoff), similar to active inter-RNC handoff procedures discussed below.

In some implementations, the central load tracker triggers the load balancing proactively. For example, when a central load tracker detects that the RNC's server cards are unevenly loaded, the central load tracking entity sends an active session transfer request to an overloaded server card along with the identity of a less-loaded server card that can serve the connection.

In some implementations, load balancing is performed across multiple RNC's in a cluster in an IP-based radio network. For example, a load tracking entity central to all RNC's in a cluster keeps tracks of the overall load of individual RNC's, by interacting with the local load trackers in each RNCs. If the central load tracking entity detects an load imbalance among the clustered RNCs, it triggers active session transfer from an overloaded RNC to a less-loaded RNC. A central tracking entity for a cluster may reside in one of the RNC's in the cluster or may be external to the RNCs in the cluster.

Load balancing and overload control mechanisms described above can be implemented in a distributed manner inside an RNC and/or inside the cluster. For example, instead of a central load tracking entity associated with a cluster, local load

tracking entities in the RNCs within the cluster can share loading information directly with each other. This way, when an RNC is experiencing overload, it can determine which RNC in the cluster to transfer an active session by, for example, contacting other local load tracking entities to determine availability of other RNC's in the cluster. It
5 can then trigger an active handoff to an available RNC. A similar strategy could be used to handle load balancing within an RNC.

When a server card fails, all connections being handled by that server card are also lost, unless there is some mechanism to save critical connection state information elsewhere in the RNC or externally in the cluster. In some
10 implementations, connection state information may be served on a system controller card of the RNC. Thus, when a server card fails, the system controller can allocate the traffic channels on the failed server card to other server cards in the system. A server that assumes the responsibility for a server card, would then establish the traffic channel legs to all sectors serving the connection, restore the A10 link to the PDSN,
15 and initialize the radio link protocol, and resume traffic channel operation.

Dormant Handoffs in Partially-Connected Radio Networks

In an IP-based radio network, a dormant AT may connect to the network via a plurality of RNCs regardless of the RN that is currently serving it. These RNC's together with the RNs they serve form a cluster. When every RN has full association
20 with every RNC, such a cluster is referred to as a mesh RNC cluster. Inside a mesh RNC cluster, a dormant AT can always maintain connectivity to its serving RNC, since the serving RNC can communicate with the AT via any one of the RNs in the cluster. This means that the serving RNC can page the AT anywhere inside the cluster and the dormant AT can send an access channel message to the serving RNC anywhere inside
25 the cluster.

When an RNs does not have an association with each RNC in a cluster, the cluster is referred to as a partially-connected cluster. In a partially-connected cluster, an AT may lose network connectivity if the RN currently serving it does not have an association with its serving RNC. This means the AT may become
30 unreachable or it may not be able to send access channel messages to its serving RNC, for example to request a new connection. To prevent this from happening, the AT's session is moved from the serving RNC to another RNC where the AT can maintain

connectivity.

In some implementations, RNCs and RNs in a cluster are subdivided into multiple mesh clusters and 1xEV-DO subnet boundaries are created between RN's that belong to different clusters. Then, whenever a dormant AT crosses the subnet
5 boundary between clusters, it sends a UATIRequest, triggering a so-called A13 inter-RNC dormant handoff. The RN who receives the UATIRequest would forward it to its primary RNC (or one of the RNC's in its cluster), and this RNC will then determine, from the UATI, the RNC currently serving the AT's session and would trigger the A13 dormant handoff. Creating subnets between clusters would thus ensure that a dormant
10 AT always remains connected to the network.

When an active AT crosses the subnet boundary between clusters, the connection can still be maintained as long as the serving RNC has the necessary association to assign traffic channels to RN's in other clusters. Otherwise, an active RNC handoff has to be performed before reaching the boundary of connectivity.

15 In some implementations, a partially-connected cluster is configured to allow an AT to move to the coverage area of an RN which does not have an association with the AT's serving RNC, and the serving RNC may not be aware that the AT is no longer reachable. For example, the current paging area of the AT, as determined by the RouteUpdateRadius of the sector where the AT last sent a RouteUpdate, may include
20 an RN that does not have an association with the serving RNC. Therefore, when the RNC receives paging data for this AT, it cannot directly page the AT via these RNs. To address this issue, the RNC maintains a paging area list for every RN with which it maintains an association. The paging area list is determined by the RouteUpdateRadius of the RN. This list not only includes the IP address of all the RN's in the paging area
25 with whom the RNC has an association, but also includes the IP address of the primary RNC for RNs with which it has no association. When the serving RNC receives paging data, it sends a Page message directly to RN's with whom it has an association and forwards the Page request for other RN's to their primary RNC, which will in turn page the AT via their RN's in the paging area. The serving RNC, when it forwards a
30 page message to another RNC, includes additional information to allow the other RNC to determine via which of its RN's it needs to send the page message.

In some implementations, the serving RNC sends a page message

directly only via those RN's for which it is the primary RNC, and for other RNs, it forwards the page message to their primary RNC, which in turn page the AT via their RN's in the paging area.

In a partially-connected cluster, an AT may send an access channel
5 message to the network when it is within the coverage area of an RN which has no association with the serving RNC. In this case, the serving RN, upon recognizing from the UATI field in the access channel packet that the packet is being served by an RNC with which it has no association, forwards the packet to its primary RNC (or some other RNC it selects based on a load balancing algorithm). The RN's primary RNC, before
10 processing the received access channel message, will trigger an A13 dormant handoff where the AT's session will be transferred from the serving RNC, which will assume the role of the Source RNC during the handoff, to the RN's primary RNC, which will assume the role of the Target RNC during the handoff. The Target RNC derives the IP address of the Source RNC from the UATI in the received access channel packet and
15 then sends an A13 request message to the Source RNC. After receiving the AT's session information from the Source RNC, the Target RNC will assign the AT a new UATI, complete the A13 handoff and then process the received access channel message. For example, if the access channel message was a ConnectionRequest message, the Target RNC will set up the traffic channel legs to the RN's involved and
20 will send a Traffic Channel Assignment message. It is important that the Target RNC complete the A13 handoff procedure in time to send the TCA message before a timer expires in the AT.

The methods described in this section can also be used in distributed IP-based radio networks where the RNC and RN functions are co-located or integrated in the
25 same base station.

Active Handoffs in Partially-Connected Radio Networks

In an IP-based radio access network, an RNC may not have an association with all RNs in the network. Therefore, a serving RNC cannot maintain a traffic channel for an access terminal (AT) indefinitely. If the AT moves into the coverage area of an RN
30 that does not have an association with the serving RNC, it will gradually lose radio connectivity to its serving RNC. Once the connection is lost, the AT will try to re-establish a new connection, initiating the activity-triggered dormant handoff procedure

described above. One way to reduce the interruption caused by lack of association is for the serving RNC, upon recognizing that the AT can better be served by sectors with whom it has no connectivity, is to close the connection, so as to prevent a potentially lengthy period before the AT decides to close the connection. Such pro-active action
5 by the serving RNC can avoid the so-called "RF dragging problem".

To prevent the connection drop described above altogether, a handoff can be performed from the serving RNC to another RNC that has association with RN's that are in the vicinity of AT's current location.

As the AT moves from the coverage area of one RN to the coverage area of
10 another RN, it sends a RouteUpdate message to its serving RNC, which the serving RNC uses to perform normal soft handoffs. Through these RouteUpdate messages, the serving RNC obtains a fairly accurate assessment of where the AT is located, and can use this information to trigger the active RNC handoff, as follows.

Each RNC may be configured with a table that shows the primary RNC of each
15 homed RN. When the RNC decides to trigger an active handoff, it can send a HandoffRequest message to the primary RNC of the serving RN to initiate the handoff. Other methods could also be devised in selecting the Target RNC for the active handoff.

Active handoff is more complex to implement than a dormant handoff, since in
20 active handoff also active call state needs to be transferred from the Source RNC to the Target RNC. Also, when a Radio Link Protocol (RLP) is in use, it becomes important to move the RLP handling from the Source RNC to the Target RNC with minimal interruption. There are alternative ways of implementing active handoff. For example, one could establish a new RLP link via the new RNC, before dropping the existing RLP
25 link via the old RNC. This allows the AT to continue to receive via the Source RNC, while the handoff is in progress. Once the handoff is completed, the RLP link to the Source RNC could be terminated and all packets could flow via the Target RNC. 3GPP2 is currently standardizing protocols to implement such active handoffs to enable interoperability between different vendor's RNC equipment.

30 The methods described in this section can also be used in distributed IP-based radio networks where the RNC and RN functions are co-located or integrated in the same base station.

Multi-Carrier Operation in IP-based Networks

Another advantage of an IP-based radio access network is in high-capacity multi-carrier (multiple frequency channel) deployments. In traditional systems with strict RNC boundaries, adding a new carrier may require splitting the RN's served by the RNC on carrier #1, so that a new set of co-located RN's on carrier #2 can be served by the same RNC. Serving co-located carriers by the same RNC is a critical requirement in traditional systems, in order to ensure that proper multi-carrier load balancing algorithms can be applied. Multi-carrier load balancing is used to distribute the traffic channel load across the available carriers equitably, so as to maximize overall user experience. With load balancing, packet data users' experience is improved, and blocking probability for voice users is minimized. As a result, when a new carrier is added, it is necessary to remove some of the RN's on the first carrier from their serving RNC in order to make room for RN's on the second carrier. This causes unnecessary service disruption.

In an IP-based radio access network, since an RNC can maintain association with a larger set of RN's, it is not necessary to remove any of the existing RN's on the first carrier from their RNC(s). Instead, one can add a 2nd RNC and have all existing RN's associate with it, while still maintaining their association with the 1st RNC. This avoids any service disruption. As more RN's are added on the 2nd carrier, these establish associations with both the 1st and 2nd RNC.

When an AT requests a new session, the serving RN can forward the session request to any one of the two RNC's based on some inter-RNC load balancing mechanism, or it may forward the session to a default RNC. This RNC will then become the serving RNC for that session, and will serve the AT as long as the AT remains within the footprint of the entire RNC cluster, regardless of which carrier the AT is monitoring. Also, at the time of traffic channel establishment, the serving RNC can assign the AT to a set of RN's that operate on any one of the available carriers. Provided the RN's supply their load information to all RNC's with whom they are associated, such traffic channel assignment can be implemented so as to balance the load across carriers. Since all RN's on all carriers may operate on the same 1xEV-DO subnet, an inter-RNC dormant handoff is not required when the AT changes the carrier it is monitoring.

Alternatively, in IP-based radio access networks with primary/secondary association, the RN's operating on the 1st carrier will use RNC #1 as their primary RNC, while RN's operating on the 2nd carrier will use RNC #2 as their primary RNC. At the same time, all RN's will have secondary association with the other RNC as well.

5 In this case, when a new session request arrives on the 1st carrier, the serving RN forwards it to its primary RNC, who then becomes the serving RNC. Alternatively, when a new session request arrives on the 2nd carrier, the serving RN forwards it to its secondary RNC, who becomes the serving RNC. Again either RNC can assign an AT to traffic channels on either one of the carriers, so as to balance the traffic load across
10 the carriers. Again this requires that the RN's supply their load information to all RNC's, both primary and secondary. The only drawback of the IP-based radio access network with primary/secondary association is that if the AT were to change the carrier it is monitoring, an inter-RNC dormant handoff has to be performed. An AT may change the carrier it is monitoring, for example, after being redirected by its serving
15 RNC. Such redirection may occur, for example, to direct an AT to a carrier that is better equipped to serve the AT. In mixed revision networks, where one carrier supports so-called Revision 0 of the 1xEV-DO standard, while a 2nd carrier supports the so-called Revision A, the network may redirect Rev A users from the Rev 0 carrier to the Rev A carrier. Such redirection may also occur at carrier boundaries, where the
20 coverage of a carrier may end, or in case of failure of an RN. If an AT loses RF link to a given RN on carrier #1, it may connect to co-located RN's on carrier #2. Again, an inter-RNC dormant handoff is required to perform such an inter-carrier handoff.

The methods described in this section can also be used in distributed IP-based radio networks where the RNC and RN functions are co-located or integrated in the
25 same base station.

Border Cell Paging

In IP-based radio networks, handoff boundaries can be greatly reduced by forming large clusters. However, formation of large clusters may not completely eliminate handoffs. A well-known problem often seen in dormant handoffs is loss of
30 paging data during dormant handoffs. In dormant inter-RNC handoffs, the AT will be monitoring only one sector at a time. As soon as the AT starts monitoring a sector that is not reachable by the serving RNC, the serving RNC can no longer page the AT. AT

reaches pageability again after its session has been transferred to a new RNC and the PDSN is now forwarding AT packets to the new RNC. In order to prevent loss of paging data, serving RNC may include in the session information that it sends to the new RNC during A13 dormant handoff any outstanding paging data. The serving
5 RNC, upon completing the A13 dormant handoff, can then page the AT. This will ensure that the AT remains page-able during the dormant inter-RNC paging procedure.

A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention, and, accordingly, other embodiments are within
10 the scope of the following claims.

1. A method comprising,
in connection with a mobile wireless network including a first and a second
radio network controller and a first and a second radio node,
establishing a first session for a first mobile access terminal on the first radio
5 network controller via the first radio node,
establishing a second session for a second mobile access terminal on the second
radio network controller via the second radio node,
establishing a first traffic channel between the first mobile access terminal and
the first radio network controller,
10 sending and receiving a first plurality of packets over the first traffic channel,
the first plurality of packets traveling between a first radio node and the first radio
network controller without passing through the second radio network controller,
maintaining the first traffic channel as the first access terminal moves from a
coverage area of the first radio node to a coverage area of the second radio node,
15 sending and receiving a second plurality of packets over the first traffic channel, the
second plurality of packets traveling between the second radio node and the first radio
network controller without passing through the first radio network controller, and
establishing a new session for the first mobile access terminal on the second
radio network controller or performing an inter-RNC dormant handoff from the first
20 radio network controller to the second radio network controller when the first access
terminal is in a dormant state after moving from a coverage area of the first radio node
to a coverage area of the second radio node.
2. The method of claim 1 further comprising,
establishing a primary association between the second radio node and the
25 second radio network controller and establishing a secondary association between the
second radio node and the first radio network controller.
3. The method of claim 2,
wherein establishing the primary association comprises passing PN offset and
IP address information from the second radio node to the second radio network
30 controller, and

wherein establishing the secondary association comprises passing PN offset and IP address information from the second radio node to the first radio network controller.

4. The method of claim 1 or 2 further comprising,
broadcasting, by the first radio node, a first subnet identifier,
5 broadcasting, by the second radio node, a second subnet identifier different from the first subnet identifier,
monitoring, by the first access terminal, the subnet identifier while in the dormant state, and
triggering the new session establishment or the dormant inter-RNC handoff
10 from the first radio network controller to the second radio network controller upon detecting a change in the subnet identifier.
5. The method of claim 4 further comprising,
individually configuring each radio node with its respective subnet identifier.
6. The method of claim 4 further comprising,
15 obtaining, by the second or a third radio node its subnet identifier from the respective radio network controller with whom it has established a primary association.
7. The method of claim 4 wherein,
triggering comprises triggering, by the first access terminal, the new session establishment or the dormant inter-RNC handoff from the first radio network controller
20 to the second radio network controller, by sending a UATI Request message.
8. The method of claim 7 further comprising,
individually configuring each radio node with its respective subnet identifier.
9. The method of claim 7 further comprising,
obtaining, by the second or a third radio node its subnet identifier from the
25 respective radio network controller with whom it has established a primary association.
10. The method of claim 4 further comprising,
establishing a primary association between a third radio node and the first radio network controller,

establishing a secondary association between the third radio node and the second radio network controller,

broadcasting, by the third radio node, a third subnet identifier different from first and second subnet identifiers associated with the first and second radio nodes,

5 the first radio network controller assigning a new UATI to the first access terminal when it is in dormant state after moving from a coverage area of the first radio node to a coverage area of the third radio node, and

triggering the new UATI assignment, by sending a UATI Request message, upon detecting a change in the subnet identifier.

10 11. The method of claim 10 further comprising,
individually configuring each radio node with its respective subnet identifier.

12. The method of claim 10 further comprising,
obtaining, by the second or a third radio node its subnet identifier from the respective radio network controller with whom it has established a primary association.

15 13. The method of claim 2 further comprising,
employing an RNC resource control agent for storing the session information for sessions being served by one or more radio network controllers.

14. The method of claim 7 further comprising,
detecting, by the RNC resource control agent, a failure of a radio network
20 controller, and
upon detection of the failure of a radio network controller, reassigning user sessions to remaining radio network controllers and passing session information to these remaining radio network controllers.

15. The method of claims 1 or 2 further comprising,
25 employing a chassis-based hardware platform with multiple server cards to implement each of the radio network controllers.

16. The method of claim 15 further comprising,
establishing an association between the second radio node and the first radio

network controller by homing on one of the server cards on the first radio network controller, and

establishing an association between the second radio node and the second radio network controller by homing on one of the server cards on the second radio network controller.

17. The method of claim 16 wherein,

establishing the association comprises passing PN offset information from the second radio node to the server card to which it is homed,

the method further comprising distributing the PN offset information from the server cards receiving the PN offset information to other server cards in the chassis-based hardware platform.

18. The method of claim 16 further comprising,

establishing a transport-layer connection for signaling between each radio node and the server cards to which it is homed in the radio network controllers.

19. The method of claim 18 wherein,

establishing the association comprises passing PN offset information from the second radio node to the server card to which it is homed,

the method further comprising distributing the PN offset information from the server cards receiving the PN offset information to other server cards in the chassis-based hardware platform.

20. The method of claim 1 wherein,

the inter-RNC handoff procedure complies with an A13 interface of a 1xEV-DO IOS specification.

21. The method of claim 1 wherein,

the radio network controllers comprise a PDSN function.

22. The method of claim 1 wherein,

the first radio network controller and the first radio node are co-located, and the second radio network controller and the second radio node are co-located.

23. A computer program product, tangibly embodied in an information carrier, and adapted to work in a mobile wireless network including a first and a second radio network controller and a first and a second radio node, the computer program product comprising instructions operable to cause data processing apparatus to:

5 establish a first session for a first mobile access terminal on the first radio network controller via the first radio node,

 establish a second session for a second mobile access terminal on the second radio network controller via the second radio node,

 establish a first traffic channel between the first mobile access terminal and the
10 first radio network controller,

 send and receive a first plurality of packets over the first traffic channel, the first plurality of packets traveling between a first radio node and the first radio network controller without passing through the second radio network controller, and

 maintain the first traffic channel as the first access terminal moves from a
15 coverage area of the first radio node to a coverage area of the second radio node, a second plurality of packets traveling between the second radio node and the first radio network controller without passing through another radio network controller.

24. The computer program product of claim 23, wherein the instructions are further operable to cause the data processing apparatus to,

20 establish a new session for the first mobile access terminal on the second radio network controller or performing an inter-RNC dormant handoff from the first radio network controller to the second radio network controller when the first access terminal is in dormant state after moving from a coverage area of the first radio node to a coverage area of the second radio node.

25 25. A mobile wireless network comprising,

 a first radio network controller,

 a second radio network controller,

 a first radio node,

 a second radio node,

30 a first mobile access terminal that is associated with a first session established

on the first radio network controller via the first radio node and a first traffic channel established with the first radio network controller, the first mobile access terminal sending and receiving a first plurality of packets over the first traffic channel, wherein the first plurality of packets travel between a first radio node and the first radio network
5 controller without passing through the second radio network controller, and
a second mobile access terminal that is associated with a second session established on the second radio network controller via the second radio node,
wherein the first traffic channel is maintained as the first access terminal moves from a coverage area of the first radio node to a coverage area of the second radio node,
10 and wherein a second plurality of packets travel between the second radio node and the first radio network controller without passing through the second radio network controller.

26. The mobile wireless network of claim 20 wherein a new session is established for the first mobile access terminal on the second radio network controller or an inter-
15 RNC dormant handoff from the first radio network controller to the second radio network controller is performed when the first access terminal is in dormant state after moving from a coverage area of the first radio node to a coverage area of the second radio node.

27. A method for exchanging digital information using mobile access terminals in a
20 wireless network, the method comprising:

transmitting packets over a first traffic channel, established between a first mobile access terminal and a first radio network controller, via a first radio node without passing through a second radio network controller;

transmitting packets over a second traffic channel, established between a second
25 mobile access terminal and the second radio network controller, via a second radio node without passing through the first radio network controller;

maintaining the first traffic channel between the first access terminal and the first radio network controller when packets are received from or transmitted to the first access terminal via the second radio node; and

30 transmitting packets received from the first access terminal to an external

network.

28. The method of claim 27 further comprising:

performing an active handoff of the first traffic channel from the first radio network controller to the second radio network controller upon detecting that the first
5 access terminal is near a third radio node.

29. The method of claim 27 further comprising:

closing the first traffic channel by the first radio network controller upon detecting that the first access terminal is near a third radio node.

30. The method of claim 27 further comprising:

10 establishing a first session for a third mobile access terminal on the first radio network controller via the first radio node;

keeping track of the approximate location of the third access terminal based on location update messages sent by the access terminal;

receiving packets at the first radio network controller for the third access
15 terminal while it is in a dormant state;

sending the second radio controller a message requesting that it page the third access terminal; and

20 paging the access terminal from the second radio network controller via a third radio node.

31. The method of claim 30 further comprising,

receiving at the third radio node a traffic channel request message from the third mobile access terminal over an access channel;

forwarding the traffic channel request message from the third radio node to the second radio network controller;

25 performing a handoff of the first session from the first radio network controller to the second radio network controller; and

after completing the handoff, establishing a traffic channel between the second radio network controller and the third access terminal.

32. A method of claim 27, wherein the first radio network controller comprises a

plurality of server cards, the method further comprising:

establishing the first traffic channel between the radio network controller and the first mobile access terminal on one of a server card selected from the plurality of server cards;

5 sending an address of the selected server card from the radio network controller to the first radio node;

 sending reverse link traffic channel packets from the first radio node to the address of the selected server card;

 sending the address of the selected server card from the first radio network
10 controller to the second radio node; and

 sending reverse link traffic channel packets from the second radio node to the address of the selected server card.

33. The method of claim 27 wherein the first radio node operates on a first frequency channel and the second radio node is co-located with the first radio node and
15 operates on a second frequency channel.

34. The method of claim 33 further comprising,

 establishing a first session for the first mobile access terminal on the first radio network controller via the first radio node operating on the first frequency channel;

 establishing a second session for the second mobile access terminal on the
20 second radio network controller via the second radio node operating on the second frequency channel; and

 establishing a third session for a third mobile access terminal on the first radio network controller via the second radio node operating on the second frequency channel.

25 35. The method of claim 33, further comprising:

 establishing a primary association between the first radio node operating on the first frequency channel and the first radio network controller;

 establishing a secondary association between the second radio node operating on the second frequency channel and the first radio network controller;

30 establishing a first session for the first mobile access terminal on the first radio

network controller via the first radio node operating on the first frequency channel; and
transferring the first session to the second radio network controller, whenever
an access terminal starts monitoring the second radio node in a dormant mode.

36. The method of claim 33 further comprising:

5 receiving a connection request from the first mobile access terminal in the first
radio network controller via the first radio node operating on the first frequency
channel; and

establishing a third traffic channel for the first access terminal on the first radio
network controller, the traffic channel flowing via the second radio node operating on
10 the second frequency channel.

37. The method of claim 36 further comprising:

establishing the third traffic channel flowing via the second radio node
operating on the second frequency channel based on actual load information received
by the first radio network controller from the first and second radio node.

15 38. The method of claim 36 further comprising:

establishing a third traffic channel flowing via the second radio node operating
on the second frequency channel based on an estimate of load on the first and second
radio n

39. The method of claim 27, wherein the first and second radio network controllers
20 each include a plurality of server cards, the method further comprising:

handling a plurality of traffic channels on a first server card in the first radio
network controller;

detecting an overload condition in the first server card;

selecting a traffic channel served on the first server card for transfer to another
25 server card; and

transferring the selected traffic channel to another server card in one of the radio
network controllers without dropping the traffic channel.

40. The method of claim 39 wherein selecting the traffic channel for transfer is
based at least partially on the amount of processing resources used by the selected

traffic channel.

41. The method of claim 39 wherein selecting the traffic channel for transfer is based at least partially on a quality-of-service requirement of the traffic on the plurality of traffic channels.

5 42. The method of claim 39 further comprising:
determining a target server card to transfer the selected traffic channel to based at least partially on load and availability of other cards in the radio network controllers.

43. The method of claim 39 wherein the first server card and target server card are both located in the same radio network controller.

10 44. The method of claim 42 wherein information about the load of server cards is provided by a centralized load tracker.

45. The method of claim 44 wherein the centralized load tracker is located in a radio network controller.

46. The method of claim 44 wherein the centralized load tracker is external to all
15 radio network controllers.

47. The method of claim 44 wherein the centralized load tracker is configured to trigger a transfer of a traffic channel from the first server card to the target server card.

48. The method of claim 41 wherein load information is obtained by the first server card directly from other server cards.

20 49. The method of claim 30, wherein the third radio node and the first radio node belong to different sub-networks.

50. The method of claim 27 further comprising:
using an Internet Protocol network to exchange data packets between the first radio network controllers and the first and second radio nodes.

25 51. The method of claim 27 wherein the first radio network controller is configured

to function as a packet data switching node.

52. The method of claim 27 wherein the first and second radio network controllers are co-located with the first and second radio nodes.

53. The method of claim 27 wherein the external network comprises an Internet
5 Protocol network.

54. A radio access network for wirelessly communicating with mobile access terminals, the radio access network comprising:

a plurality of radio nodes interconnected with a plurality of radio network controllers using a network, wherein each said radio node can address each said radio
10 network controller and each said radio network controller can address each said radio node; and

an interface for exchanging packets between the radio access network and an external network.

55. The radio access network of claim 54 wherein the network interconnecting the
15 radio nodes and radio network controllers comprises an Internet Protocol network.

56. The radio access network of claim 54 wherein each radio network controller is configured to maintain a traffic channel regardless of which radio node the traffic channel is flowing.

57. The radio access network of claim 56 wherein the radio network controllers
20 maintain the traffic channel by routing packets to and from the access terminal via any one of the radio nodes.

58. The radio access network of claim 56 wherein the plurality of radio nodes and the plurality of radio network controllers are associated with a common sub-network.

59. The radio access network of claim 54 wherein each of the radio nodes is
25 associated with a primary radio network controller selected from the plurality of radio network controllers.

60. The radio access network of claim 54 wherein each of the plurality of radio network controllers are enabled to send a paging message to an access terminal via any of the radio nodes.

61. The radio access network of claim 54 wherein each radio network controller
5 comprises:

a plurality of server cards each connected to the network and addressable by each of the plurality of radio nodes, wherein each radio network controller is configured to establish traffic channels with access terminals on each of the plurality of server cards.

10 62. The radio access network of claim 61 wherein the radio network controller is configured to provide an address of a server card on which a traffic channel is established to one or more radio nodes.

63. The radio access network of claim 54 wherein the plurality of radio nodes comprises:

15 a first radio node configured to operate on a first frequency channel; and
a second radio node configured to operate on a second frequency channel,
wherein the first and second radio nodes are co-located.

64. The radio access network of claim 61 wherein each radio network controller is configured to detect an overload condition in one of its plurality of server cards and
20 select one or more traffic channels handled by an overloaded card for transfer to another card.

65. The radio access network of claim 54 wherein the interface comprises a packet data switching node.

66. The radio access network of claim 54 wherein the interface is part of a radio
25 network controller.




FIG. 1




FIG. 2

3/4




FIG. 3

4/4




FIG. 4

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
1 December 2005 (01.12.2005)

PCT

(10) International Publication Number
WO 2005/115026 A3

(51) International Patent Classification:

H04Q 7/20 (2006.01) *H04L* 12/46 (2006.01)
H04Q 7/00 (2006.01) *H04B* 7/12 (2006.01)

(21) International Application Number:

PCT/US2005/017385

(22) International Filing Date: 17 May 2005 (17.05.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

10/848,597 18 May 2004 (18.05.2004) US
11/037,896 18 January 2005 (18.01.2005) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:

US 11/037,896 (CON)
Filed on 18 January 2005 (18.01.2005)
US 10/848,597 (CIP)
Filed on 18 May 2004 (18.05.2004)

(71) Applicant (for all designated States except US): **AIR-VANA, INC.** [US/US]; 25 Industrial Avenue, Chelmsford, MA 01720 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **EYUBOGLU, Vedat** [US/US]; 150 Jennie Dugan Road, Concord, MA

01742 (US). **KIM, Woojune** [KR/US]; 53 Lake St #2, Arlington, MA 02474 (US). **BARABELL, Arthur, J.** [US/US]; 11 Hayden Circle, Sudbury, MA 01776-2098 (US). **CHERIAN, Sanjay** [US/US]; 6 Maxwell Drive, Brookline, MA 03033 (US).


(74) Agent: **FEIGENBAUM, David, L.**; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US (patent), UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: RADIO NETWORK CONTROL



(57) Abstract: A first session is established for a first mobile access terminal (AT) on a first radio network controller (RNC) via a first radio node (RN). A first traffic channel is established between the first mobile access terminal (AT) and the first radio network controller (RNC). A first plurality of packets are sent and received over the first traffic channel. The first plurality of packets travel between a first radio node (RN) and the first radio network controller (RNC) without passing through a second radio network controller (RNC). The first traffic channel is maintained as the first access terminal moves from a coverage area of the first radio node (AT) to a coverage area of a second radio node (RN). A second plurality of packets travel between the second radio node (RN) and the first radio network controller (RNC) without passing through another radio network controller.

WO 2005/115026 A3



Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(88) Date of publication of the international search report:

4 January 2007

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US05/17385

A. CLASSIFICATION OF SUBJECT MATTER

IPC: **H04Q 7/00**(2006.01);**7/20**(2006.01);**H04L 12/46**(2006.01)
H04B 7/212(2006.01)

USPC: 370/328,331;455/432.1,436,442

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/328-331, 337, 341, 443; 455/432.1, 434,436, 440, 442, 445, 450

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
Please See Continuation Sheet

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6,408,182 B1 (DAVIDSON et al.) 18 June 2002, column 3-5.	54-66
Y	WO 98/08353 A2 (HOKKANEN) 26 February 1998, page 7-14.	1-14, 20-29,33-38 and 50-53
Y	US 2002/0193110 A1 (JULKA et al.) 19 December 2002, paragraphs 0029, 0039, 0077.	1-14, 20-29, 33-38 and 50-53
Y	US 6,580,699 B1 (MANNING et al) 17 June 2003, column 1-2.	2
Y	US 2002/0145990 A1 (SAYEEDI) 10 October 2002, paragraph 0031.	4-6
Y	US 2002/0067707 A1 (MORALES et al.) 06 June 2002, Fig. 1 and paragraph 0056	7-12, 14 and 20-22
X	US 6,366,961 B1 (SUBBIAH et al.) 02 April 2002, column 4, lines 54-67.	54-66

☒ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T"

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X"

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y"

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&"

document member of the same patent family

Date of the actual completion of the international search

08 September 2006 (08.09.2006)

Date of mailing of the international search report

26 OCT 2006

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

Facsimile No. (571) 273-3201

Authorized officer

Saba Tsegaye

Telephone No. 571-272-2600

Rugenia Tsegaye

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US05/17385

C. (Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X, P	US 6,738,625 B1 (OOM et al) 18 May 2004, enire document.	1, 2, 4-6, 13, 15-19, 22-49 and 52

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US05/17385

Continuation of B. FIELDS SEARCHED Item 3:

EAST

search terms: mobile, wireless, radio network controller, base station, RNC, handoff